The Secret Lives of Assumptions: Developing and Refining Assumption Personas for Secure System Design

Shamal Faily and Ivan Fléchaïs

Oxford University Computing Laboratory
Wolfson Building, Parks Road, Oxford OX1 3QD UK
{shamal.faily,ivan.flechaïs}@comlab.ox.ac.uk

Abstract. Personas are useful for obtaining an empirically grounded understanding of a secure system’s user population, its contexts of use, and possible vulnerabilities and threats endangering it. Often, however, personas need to be partly derived from assumptions; these may be embedded in a variety of different representations. Assumption Personas have been proposed as boundary objects for articulating assumptions about a user population, but no methods or tools currently exist for developing and refining these within the context of secure and usable design. This paper presents an approach for developing and refining assumption personas before and during the design of secure systems. We present a model for structuring the contribution of assumptions to assumption personas, together with a process for developing assumption personas founded on this model. We also present some preliminary results based on an application of this approach in a recent case study.

1 Introduction

Personas are useful for obtaining a grounded understanding of a system’s contexts of use, and communicating that understanding within a design team. Recent work on applying personas to help elicit and specify secure system requirements found that the data and analysis from which personas are derived also help identify threats and vulnerabilities [10]. Although adherents of personas argue that these should be primarily derived from real-world observations [7,14], the necessary resources for eliciting and analysing such data may not always be available. In these cases, it is necessary to rely on second-hand data about users and their contexts, much of which might be derived from assumptions.

Many usability professionals are familiar with analysing assumption-based usage data, but this may not be the case for software engineers. Engineers are usually employed for their technical expertise and domain knowledge; we cannot reasonably expect them to have a working knowledge of usability design techniques as well. They do, however, have tacit knowledge of the problem domain and a sensitivity to the values at play within its contexts of use. The challenge is to not only trace assumptions made about personas to their source, but to

explicate the claims these assumptions represent. By doing so, we also explicate tacit knowledge about users and their contexts. Like data directly elicited from real-world observations, this data also suggests hitherto unknown threats and vulnerabilities related to a system.

Techniques from Design Rationale research are useful for tracking the refinement of assumptions to architectural components and software. Such techniques may also be useful for tracking the same assumptions to less refined concepts used in security analysis. Security design has the same needs for discharging potential ambiguity grounded in assumptions; these may be sources of attack vectors if the vulnerabilities they expose are exploited. In this paper, we present an approach for developing assumption personas for secure system design, and describe how this approach can be embedded into an existing design process and associated tool-support. In section 2, we briefly introduce personas and describe the related work motivating our approach. In section 3 we present an overview of our approach, and in section 4 we report on some preliminary findings which arose when applying this approach in a recent case-study.

2 Related Work

2.1 Personas and Assumption Personas

Personas are behavioural specifications of archetypical users. These were introduced by Cooper [6] to deal with programmer biases arising from the word *user*. These biases lead to programmers bending and stretching assumptions about users to meet their own expectations; Cooper called this phenomena *designing for the elastic user*. Personas are now a mainstay in User-Centered Design, with articles, book-chapters, and even a book [14] devoted to developing and applying them in practice. Personas have also been applied to Requirements Engineering, an area of intersection between HCI and Software Engineering [4].

Accepting that data-driven personas are an ideal rather than a norm, Pruitt & Adlin [14] proposed *Assumption Personas*: persona sketches created to articulate existing assumptions about an organisation's user population. These personas are grounded in assumptions contributors hold about users, and the context of investigation. These assumptions may be derived from interpreted or misinterpreted experiences, and coloured by individual and organisational values. Assumption Personas help people see the value of personas in design, and how different assumptions shape them. As a result, when exposed, they can guide subsequent analysis or data collection for data-driven personas.

Personas are not, however, without their critics. Chapman & Milham argue that, as fictional archetypes, personas are difficult to verify as there is no way to falsify them [5]. They further argue that questions remain about how personas should be reconciled with other information, understanding what data underpins their characteristics, and what happens when different interpretations are made from the same persona.

2.2 Integrating Personas with Secure Software Engineering

Chapman and Milham’s criticism about the stand-alone nature of personas can be addressed by integrating them into the software engineering process. This has been the subject of our recent work on the IRIS (Integrating Requirements and Information Security) framework, which integrates usability into the design of secure software systems [8]. As part of this work, a meta-model for usable secure requirements engineering was devised, which integrates the persona with other concepts in usability, security, and software engineering. From this model, we have developed CAIRIS (Computer Aided Integration of Requirements and Information Security): a tool for managing information about personas and other design elements, and evaluating the effect to security and usability of different design decisions [1]. CAIRIS manages requirements, task, and risk data, and automatically generates different types of visual model to represent the ongoing analysis. We demonstrate this approach in [9] by illustrating how categorical information about a task performed by a pre-defined persona is associated with the results of risk analysis, and how the usability of this task can be visually represented before and after a related risk is mitigated.

In [10], we presented a process for developing personas for secure systems; this is based on collecting and analysing empirical data from qualitative and contextual interviews. The personas derived from analysing this data were validated and further refined in participatory requirements and risk analysis workshops. We also found that empirical data used to derive personas could be re-used for other analysis.

Even though personas may be grounded in empirical data, the quandary about the validity of personas remains. It may be possible to verify the quality of the empirical-data or the robustness of the methodology to develop them, but we cannot easily falsify the representativeness of personas. The vision of the system may be tentative enough that what may have been valid working assumptions at the beginning of the persona development process may be invalid by the time the personas are presented to project stakeholders. It is, therefore, useful to understand how characteristics about personas track back to their assumptions, and why.

2.3 Toulmin’s Model of Argumentation

Codifying the rationale underpinning assumption personas guides analysis and decision making, but the *rationale capture problem*, characterised by the reluctance of those involved in design activities to record their rationale, cannot be ignored [3]. Although the Design Rationale community has proposed several different approaches for building rationale capture into the design process, the Security and Requirements Engineering community has taken a recent interest in capturing rationale using the vehicle of informal argumentation. These approaches are founded on Toulmin’s Argumentation Model: a logical structure for reasoning about the validity of arguments [15], the elements of which are defined in table 1.

Table 1. Elements of Toulmin’s Argumentation Model

Element	Description
Claim	A proposition representing a claim being made in an argument.
Grounds	One or more propositions acting as evidence justifying the Claim.
Warrant	One or more rules of inference describing how the Grounds contribute to the Claim.
Backing	The knowledge establishing the Grounds for believing the Warrant.
Modal Qualifier	A phrase qualifying the degree of certainty in the argument for the Claim.
Rebuttal	One or more propositions challenging the validity of the Claim.

Alexander & Beus-Dukic describe a number of simple rationale models for Requirements Engineering based on this structure [2]. From a security standpoint, Haley et al. have proposed using Toulmin’s model to support arguments for security requirements [11]. In their approach, an argument for a system satisfying its security requirements is presented for analysis. Each proposition within this argument is treated as a Claim, and argued accordingly. Rebuttals represent *Trust Assumptions*; these can be countered as part of another security argument, or examined in subsequent threat modelling activities.

3 Approach

We have chosen to embrace, rather than ignore, the contribution assumptions make to assumption persona design. We propose a novel approach to structuring the contribution of assumptions to persona specifications, and integrating this conceptual structure into an existing approach for secure systems design.

3.1 Developing Assumption Personas

Personas are usually represented as a narrative describing the behaviour of an archetypical user. Authoring these narratives remains a creative exercise, but we propose augmenting these by structuring the assumption data contributing to them. We have aligned this structure to Toulmin’s Model of Argumentation, introduced in section 2.3. Adopting this approach allows us to treat assumptions directly contributing to part of the narrative as a Claim. The task of justifying this Claim both strengthens the foundations of the persona, and guides subsequent elicitation and analysis activities. These Claims are represented conceptually using one of more *Characteristics*; these are propositions about a specific aspect of a persona’s behaviour. Characteristics are categorised according to one of the behavioural variable types defined by IRIS personas; these are based

Fig. 1. Conceptual model of assumption persona data (left), and Toulmin model visualisation based on an individual characteristic (right)

on the behavioural variable types proposed by Cooper [7]: activities, attitudes, aptitudes, motivations, and skills. Also associated with a Characteristic is a qualifying phrase representing the strength of belief in the Characteristic; this qualifying phrase aligns with the Modal Qualifier in Toulmin’s model.

Persona Characteristics originate from one of two sources. The first source is some form of *Artifact*: a document related to the problem domain or the system being specified, such as a specification, or a transcript from an interview or design workshop. The second is a design *Concept*: an instance of an object defined within the work-in-progress IRIS analysis, such as a description for an asset, a goal or requirement, or even another persona. Because an individual source may give rise to multiple Characteristics of the same or different behavioural categories, a *Reference* is associated with a given source and Characteristic. The contents of a Reference will depend on the source type. In the case of an Artifact, a reference contains information tying an attributable piece of information or comment to a source document or verbal comment, e.g. page number, document version, or person. In the case of a Concept, a Reference contains the name and type of the contributing concept. In both cases, the Reference will contain as much textual attribution information as necessary to justify the persona’s Characteristic. The name of the Reference object is a synoptic proposition of this attribution information. With regards to Toulmin’s model, References align to either Grounds, Warrants, or Rebuttals. Where a Reference represents a Warrant, the corresponding Artifact or Concept acts as the Warrant’s Backing.

The meta-model in figure 1 (left) summarises these concepts and their relationships. The stereotypes adjacent to each class represent the corresponding concept name from Toulmin’s model.

3.2 Applying and Refining the Assumption Personas

Before assumption personas are used, they are presented to a workshop or focus group containing representative system stakeholders. Following this workshop, the remaining steps of the process are carried out in the context of smaller design sessions, as described by [10]. These sessions entail requirements and risk analysis activities, where, rather than referring to users, personas are used in their place. In both the workshop and design sessions, new assumptions about personas may be identified, or existing assumptions challenged. Armed with the proposed meta-model, tool-support can be developed to elicit the structural elements of the assumption persona argumentation model. Aside from guiding and structuring the elicitation of assumption data, the structured argument of Characteristics can be cross-checked with the persona narrative. If it becomes difficult to write a believable narrative based on the Characteristics identified, then these need to be re-evaluated.

We modified the CAIRIS tool introduced in section 2.2 to illustrate how tool-support can take advantage of this approach. As well as allowing Characteristics associated with a persona to be quickly reviewed against the narrative, we found that Characteristics could be quickly created or modified when assumptions are introduced or challenged during design sessions. Structuring the data according to the meta-model also facilitates the automatic generation of visual Toulmin models for persona Characteristics. An example of such a model for a specific Characteristic is provided in figure 1 (right).

Unsubstantiated Claims and Rebuttals are also an additional source of risk analysis information. In the case of the latter, obstacles – conditions representing undesired behaviour preventing an associated goal from being achieved [12] – can be elicited from these, and its placement guided by the related Characteristic negated by the Rebuttal. This placement guidance is possible because a persona invariably participates in tasks operationalised by one or more goals or requirements.

4 Preliminary Results

We used this approach to help specify requirements for an online portal for a medical research project. The nature of this project was such that eliciting empirical data from representative users during the study was impossible. During the course of the project, an assumption persona – Alex – was developed to embody the assumptions held by the project team about the researchers expected to use the portal. The assumptions underpinning this persona were initially derived from a high-level requirements specification document developed by a different team within the same project; as such, Alex represented the assumptions that team had made about the expected user population. After developing this persona, a half-day workshop was held with the complete project team to agree the scope for a subsequent requirements and risk analysis of the portal. During this workshop, Alex was presented to the team. The team both agreed and disagreed with the characteristics of Alex. Where there was disagreement,

the structured nature of the assumption data was used to track the questionable characteristic to its originating source, which was discussed in more detail within the team. Following the workshop, a number of new assumptions were elicited, which formed the basis of new characteristics about Alex.

After the workshop, three 2-hour design sessions were held with team members to carry out requirements and risk analysis relating to two specific tasks carried out by Alex. As part of this analysis, scenarios were developed describing how Alex would carry out these tasks with the aid of the portal. During these sessions, Alex's characteristics evolved; by the end of the 3rd session, 23 different Characteristics about Alex had been captured. Some of these were modifications to assumptions captured in the initial stages of persona development, but several were derived from assumptions which surfaced while eliciting other concepts, such as tasks and goals. In all cases, these characteristics were justified by Grounds, and in many cases, a Warrant and Backing were also elicited.

Haley & Nuseibeh [13] observed that experts provide essential domain knowledge about the subtleties of threats, but non-experts ask journalist questions challenging implicit assumptions assumed by the domain expert. Our preliminary results during the design sessions concur with this observation. When the tasks carried out by one of the personas was modelled during one session, one non-expert participant raised pertinent points about implicit assumptions in the task description; these were not accounted for by the personas, and led to the rebuttal of one Characteristic.

Although identifying Grounds for Characteristics was found to be straightforward, identifying Warrants provided to be more difficult. In particular, we found that, prior to their initial validation, many of the Characteristics were based exclusively on Grounds, rather than Warrants as well. As such, value judgements about the source data and the context were directly reflected in these Characteristics. Although the initial workshop surfaced a number of these issues, it was usually not until the personas were directly used to model tasks in design sessions that many invalid Characteristics were identified. Applying the personas within a specific context did, however, help identify missing inferential data, or guide the refactoring of the argumentation structure for affected Characteristics.

5 Conclusion

Personas are a mainstay in User-Centered Design, yet there is a dearth of guidance on how to build and refine these from assumptions, as opposed to empirical data. We believe this guidance, and corresponding tool-support, may contribute to a wider adoption of personas in secure software engineering, and a better understanding of how to use these in a secure software engineering context. This paper makes three contributions towards these ends. First, we have presented a model for structuring the assumptions contributing to personas; to help guide subsequent analysis, this model has been aligned these with Toulmin's Model of Argumentation. Second, we have illustrated how tool support reifies this structured model, and guides subsequent risk analysis. Finally, we have reported some

of the preliminary results validating our approach in a recent case study. A more detailed report of this study will appear as a future publication.

Acknowledgements

The research described in this paper was funded by EPSRC CASE Studentship R07437/CN001. We are very grateful to Qinetiq Ltd for their sponsorship of this work.

References

1. CAIRIS web site, <http://www.comlab.ox.ac.uk/cairis>
2. Alexander, I., Beus-Dukic, L.: *Discovering requirements: how to specify products and services*. Wiley, Chichester (2009)
3. Burge, J.E., Carroll, J.M., McCall, R., Mistrik, I.: *Rationale-Based Software Engineering*. Springer, Heidelberg (2008)
4. Castro, J., Acua, S., Juristo, N.: Integrating the personas technique into the requirements analysis activity. In: *Mexican International Conference on Computer Science, ENC 2008*, pp. 104–112 (October 2008)
5. Chapman, C.N., Milham, R.P.: The persona's new clothes: Methodological and practical arguments against a popular method. In: *Proceedings of the Human Factors and Ergonomics Society 50th Annual Meeting* (2006)
6. Cooper, A.: *The Inmates Are Running the Asylum: Why High Tech Products Drive Us Crazy and How to Restore the Sanity*, 2nd edn. Pearson Higher Education (1999)
7. Cooper, A., Reimann, R., Cronin, D.: *About Face 3: The Essentials of Interaction Design*. Wiley, Chichester (2007)
8. Faily, S., Fléchaïs, I.: A Meta-Model for Usable Secure Requirements Engineering. In: *ICSE Workshop on Software Engineering for Secure Systems, SESS 2010*, pp. 126–135. IEEE Computer Society Press, Los Alamitos (May 2010)
9. Faily, S., Fléchaïs, I.: Analysing and Visualising Security and Usability in IRIS. In: *Fifth International Conference on Availability, Reliability and Security, ARES 2010* (2010)
10. Faily, S., Fléchaïs, I.: Barry is not the weakest link: Eliciting Secure System Requirements with Personas. In: *Proceedings of the 2010 British Computer Society Conference on Human-Computer Interaction BCS HCI 2010* (to appear, 2010)
11. Haley, C.B., Laney, R., Moffett, J.D., Nuseibeh, B.: Arguing satisfaction of security requirements. In: Mouratidis, H., Giorgini, P. (eds.) *Integrating Security and Software Engineering*, ch. 2, pp. 16–43. Idea Group, USA (2007)
12. van Lamsweerde, A., Letier, E.: Handling obstacles in goal-oriented requirements engineering. *IEEE Transactions on Software Engineering* 26(10), 978–1005 (2000)
13. Nuseibeh, B., Haley, C., Foster, C.: Securing the skies: In requirements we trust. *Computer* 42(9), 64–72 (2009)
14. Pruitt, J., Adlin, T.: *The persona lifecycle: keeping people in mind throughout product design*. Elsevier, Amsterdam (2006)
15. Toulmin, S.: *The uses of argument*. Cambridge University Press, Cambridge (2003) (updated edn.)